arXiv:1702.07778v1 [stat.ME] 24 Feb 2017

A Note on Nonlocal Prior Method in High Dimensional Setting

Yuanyuan Bian [*], *and Ho-Hsiang Wu* [†]

Abstract

We propose a new class of nonlocal prior to improve the performance of variable selection in high dimensional setting. We prove our new prior possesses the robustness to hyper parameter settings and is able to detect smaller decreasing signals.

Keywords: Nonlocal Prior, Bayesian Variable Selection, High-dimensional Data

1 Introduction

In this note we consider variable selection problem in high dimensional generalized linear models (GLMs). Let the data consists of a sequence of $\{\mathbf{Y}_n, \mathbf{X}_n\}$, where $\mathbf{Y}_n$, the responses, collect n independent observations $y_i, i = 1, ..., n$, and $\mathbf{X}_n$, the regressors, forms $n \times p$ matrix. Given $\mathbf{X}_n$, each y_i is conditionally independent and assumed to follow a distribution from the natural exponential family

$$f(y_i|\theta_i) = c(y_i) \exp\left[y_i\theta_i - b(\theta_i)\right], i = 1, , n, \tag{1}$$

with natural parameter $\boldsymbol{\theta}_n = [\theta_1, ..., \theta_n]'$. The mean function $\boldsymbol{\mu} = E(\mathbf{Y}_n|\mathbf{X}_n) = b'(\boldsymbol{\theta}_n)$. The matrix $\mathbf{X}_n$ influences $\mathbf{Y}_n$ in the form of $\boldsymbol{\theta}_n = \mathbf{X}_n\boldsymbol{\beta}$, where $\mathbf{X}_n$ can be further expressed as $[\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, ..., \mathbf{x}^{(p)}]$, a collection of regressors $\mathbf{x}^{(j)}$, for $j = 1, ..., p$, and $\boldsymbol{\beta} = [\beta_1, ..., \beta_p]'$ denotes a $p \times 1$ regression coefficient vector. This class of models include regression models whose responses are binomial, Poisson, and Gaussian with known variance.

Now suppose the joint distribution of $\mathbf{Y}_n$ is determined by a true parameter vector $\boldsymbol{\beta}_0 \in \mathcal{R}^p$ supported on a small set $\mathcal{J}_0$, such that $\mathcal{J}_0 \subset \{1, ..., p\}$ and $\beta_{0j} \neq 0$ if and only if $j \in \mathcal{J}_0$. Our interest here is to identify such set $\mathcal{J}_0$ that gives the most parsimonious true model. We are concerned with the asymptotic setting in which both n and p increase, particularly with p being of substantial size with regard to n. Specifically, we assume that

[*]University of Missouri, Columbia, MO, USA .

[†]National Cancer Institute, Bethesda, MD, USA .

(A1) $\log(p) = O(n^{\omega})$ for some $\omega \in (0, \frac{1}{3}]$,

which allows for subexponential growth of p with respect to n. The variable selection problem under such assumption is challenging and recently draws great attention, see Chen and Chen (2012); Liang et al. (2013) among others. Notably, many approaches resort to the sparsity assumption, by considering a priori bound, q, on the size of the models to be considered. By imposing such bound q, the total number of candidate models scales down dramatically from 2^p to p^q. Here we also assume the sparsity such that $|\mathcal{J}_0| \leq q$ and

(A2) $q \log(p) = O(n^{\psi})$ for some $\psi \in [0, \frac{1-\omega}{3})$,

to allow q to grow with n but at a rather low rate.

We treat the variable selection problem with nonlocal prior method, a Bayesian variable selection framework that is first introduced by Johnson and Rossell (2010). With the assignment of a nonlocal prior to induce probabilistic separation between the considered models, the nonlocal prior method has been proved to provide stronger parsimony than many of its competitors. Johnson and Rossell (2012) modified the nonlocal prior method for Gaussian linear model, and proved its consistency in $p = O(n)$ setting. Recently, Shin et al. (2017) undertook the study in $p = O(e^n)$ setting, and not only established the consistency results (with other conditions), but also demonstrated several advantages of their nonlocal prior method against other current-state-of-art methods, such as g-prior method and penalized likelihood approaches.

Despite the recent advances of nonlocal prior method, one aspect, i.e., the specification of hyper parameters of the prior, remains an open research interest. To this end, we go back to Shin's nonlocal prior method with a new class of nonlocal prior, and set out to demonstrate its potential advantages including robustness to hyper parameter settings and the ability in detecting small decreasing signal.

2 Main Results

2.1 Set up and notation

Recall the problem of interest here is to recover the support of $\boldsymbol{\beta}_0$, that is, the set

$$\mathcal{J}_0 : \{j \in [p] : \beta_{0j} \neq 0\},$$

where $|\mathcal{J}_0| \leq q$ and $[p] := \{1, ..., p\}$. We first introduce some notations for submodel. Let $\mathcal{J}$, such that $\mathcal{J} \subset [p]$, index a generic submodel consisting of a subset of p covariates. Let $\boldsymbol{\beta}_{\mathcal{J}}$

denote the corresponding length $|\mathcal{J}|$ regression coefficient vector of this submodel. Let $\hat{\boldsymbol{\beta}}_{\mathcal{J}}$ denote the MLE of $\boldsymbol{\beta}_{\mathcal{J}}$, and $\hat{\boldsymbol{\beta}}_{\mathcal{J},\text{PM}}$ denote the posterior mode of $\boldsymbol{\beta}_{\mathcal{J}}$. We simply write $\mathbb{R}^{|\mathcal{J}|}$ for the parameter space comprising all $\boldsymbol{\beta}$ with length $|\mathcal{J}|$. We further denote the log-likelihood function by $\ell(\cdot)$, the score function by $\mathcal{S}(\cdot)$, the negative Hessian of the log-likelihood function by $\mathcal{H}(\cdot)$, and the negative Hessian of the log-posterior density function by $\mathcal{H}^*(\cdot)$.

We now describe the nonlocal prior method in Shin et al. (2017). Given a generic model $\mathcal{J}$, they first assigned on the regression coefficients $\boldsymbol{\beta}_{\mathcal{J}}$ the product inverse moment prior (piMOM, Johnson and Rossell 2012),

$$\pi(\boldsymbol{\beta}_{\mathcal{J}}|r,\tau) = \frac{\tau^{r|\mathcal{J}|/2}}{\Gamma(r/2)^{|\mathcal{J}|}} \prod_{j=1}^{|\mathcal{J}|} |\beta_{\mathcal{J}j}|^{-(r+1)} \exp(-\tau/\beta_{\mathcal{J}j}^2),$$

and then carried out the marginal likelihood $\mathcal{M}_{\mathcal{J}}$, using the standard Laplace approximation. To incorporate the prior believe on sparsity, they assumed on model $\mathcal{J}$ a uniform prior $\pi_{\mathcal{J}} \propto \mathbf{I}(|\mathcal{J}| \leq q)$, so that the model space is restricted to models whose size is less than or equal to the upper bound q. Note $\mathbf{I}(\cdot)$ above denotes the indicator function. Putting all these together, Shin et al. (2017) then computed the model posterior probability for every model in consideration, and then used this quantity as criterion to identify $\mathcal{J}_0$ with the highest posterior model.

Following Shin et al. (2017), we find that the model posterior probability of the true model $\mathcal{J}_0$ can be expressed in the form of

$$p(\mathcal{J}_0|\mathbf{Y}_n) = \left[1 + \sum_{\mathcal{A}} \frac{\mathcal{M}_{\mathcal{J}}}{\mathcal{M}_{\mathcal{J}_0}} + \sum_{\mathcal{B}} \frac{\mathcal{M}_{\mathcal{J}}}{\mathcal{M}_{\mathcal{J}_0}}\right]^{-1}, \tag{2}$$

where $\mathcal{A}$ denotes the set $\{\mathcal{J} : \mathcal{J}_0 \subset \mathcal{J}, |\mathcal{J}| \leq q\}$ that collects nested models, and $\mathcal{B}$ denotes the set $\{\mathcal{J} : \mathcal{J}_0 \not\subseteq \mathcal{J}, |\mathcal{J}| \leq q\}$ that collects non-nested models. Here it is clear to see that the last two summations play crucial roles, since if they go to zero as the sample size n increases, then the method achieves the Bayesian variable selection consistency (Bayarri et al., 2012). In the literature, it is recognized that the most appealing contribution of nonlocal prior method is the improvement, when compared with other Bayesian methods that employ local priors, of the variable selection within the set $\mathcal{A}$. The convergence rate (toward zero) of the ratio of $\mathcal{M}_{\mathcal{J}}$ versus $\mathcal{M}_{\mathcal{J}_0}$ within $\mathcal{A}$ is sub-exponential with Shin's nonlocal prior method, but is polynomial with local prior method.

To see this, we first introduce below two assumptions on the Hessian $\mathcal{H}(\cdot)$:

(B1) There exist constants c_L and c_U such that for all $|\mathcal{J}| \leq q$ and all $\boldsymbol{\beta} \in \mathbb{R}^{|\mathcal{J}|}$, the negative Hessian function, $\mathcal{H}(\boldsymbol{\beta})$, is properly bounded as $c_L \mathbf{I}_{|\mathcal{J}|} \preceq n^{-1}\mathcal{H}(\boldsymbol{\beta}) \preceq c_U \mathbf{I}_{|\mathcal{J}|}$, where

the notation "$\preceq$" refers to the ordering with $\mathbf{A} \preceq \mathbf{B}$ whenever $\mathbf{A} - \mathbf{B}$ is positive semidefinite.

(B2) There is a constant c_D such that $n^{-1}||\mathcal{H}(\boldsymbol{\beta}) - \mathcal{H}(\boldsymbol{\beta}^*)||_S \leq c_D \cdot ||\boldsymbol{\beta} - \boldsymbol{\beta}^*||_2$ for all $|\mathcal{J}| \leq q$ and all $\boldsymbol{\beta}, \boldsymbol{\beta}^* \in \mathbb{R}^{|\mathcal{J}|}$, where $||\cdot||_S$ is the spectral norm of a matrix.

These two assumptions are concerned with the asymptotic identifiability. In other words, they ensure that with large sample size, the true model is always properly bounded away from the wrong model. Next, we consider a proposition to see the asymptotic behavior of $\sum_{\mathcal{A}} \mathcal{M}_{\mathcal{J}}/\mathcal{M}_{\mathcal{J}_0}$ in (2)

Proposition 1. *Suppose* piMOM *is assumed. Fix* r, ϵ, *and* $\nu > 0$. *If regularity assumptions* (A1), (A2), (B1) *and* (B2) *are satisfied, then for all* $\mathcal{J}$ *such that* $\mathcal{J}_0 \subset \mathcal{J}, |\mathcal{J}| \leq q$, *we have*

$$\log\left(\frac{\mathcal{M}_{\mathcal{J}}}{\mathcal{M}_{\mathcal{J}_0}}\right) \asymp (1+\epsilon)\log p^{(\nu+|\mathcal{J}|-|\mathcal{J}_0|)} - \left(\sum_{j\in\mathcal{J}} \tau\hat{\beta}^{-2}_{\mathcal{J}i,PM} - \sum_{j\in\mathcal{J}_0} \tau\hat{\beta}^{-2}_{\mathcal{J}_0 i,PM}\right) \tag{3}$$

where the notation "$\asymp$*" refers to the asymptotic equivalence.*

The proof is straightforward and thus omitted. In the right hand side of (3), the first term comes from the log-likelihood ratio, while the second term (the two summations within the bracket) comes from the log-prior ratio. The convergence rate of the log-likelihood ratio is established in the Theorem 2.2 of Barber et al. (2015) and holds here since they assume more general regularity conditions (at cost of more tedious technical work). From Proposition 1, we see two substantial factors determine the convergence rate of the log-prior ratio, one is the exponential kernel of piMOM, and another is the convergence rate (toward zero) of the posterior mode $\hat{\boldsymbol{\beta}}_{\mathcal{J}_0,\text{PM}}$. With additional regularity conditions and properly chosen τ, we can show that the convergence rate of the log-prior ratio is of order $\mathcal{O}(n^c)$ for some $c > 0$ such that the summation $\sum_{\mathcal{A}} \mathcal{M}_{\mathcal{J}}/\mathcal{M}_{\mathcal{J}_0}$ in (2) converges toward zero exponentially fast.

Note although Proposition 1 elucidates the advantage of using nonlocal prior method in an asymptotic sense, it does not provide practical guide on setting hyper parameter τ since the proportionality is unknown.

2.2 A new class of nonlocal prior and asymptotic results

For practical purpose, Nikooienejad et al. (2016) have proposed for setting τ a heuristic procedure that is lately shown by Shin et al. (2017) to work well with piMOM in the high-dimensional setting. Below we propose a new class of robust nonlocal prior and indicate why we think it may be of use with Nikooienejad's procedure.

Proposition 2. *Suppose we replace* τ *in* piMOM *with* τ_j *and further assume each* τ_j *follows an inverse-gamma distribution with shape* $(r+1)/2$ *and scale* λ*, then we have*

$$\begin{aligned}\pi(\boldsymbol{\beta}_{\mathcal{J}}|r,\lambda) &= \int ... \int \pi(\boldsymbol{\beta}_{\mathcal{J}}|r,\tau_j)\pi(\tau_j|\lambda)d\tau_1...d\tau_{|\mathcal{J}|} \\ &= \prod_{j=1}^{|\mathcal{J}|} \frac{|\beta_{\mathcal{J}j}|^{-(r+1)}\lambda^{(r+1)/2}\sqrt{\pi}}{2\Gamma(\frac{r+1}{2})\Gamma(\frac{r}{2})\sqrt{\lambda}} \exp\left(-2\sqrt{\frac{\lambda}{\beta^2_{\mathcal{J}j}}}\right).\end{aligned} \tag{4}$$

Remark 1. We call this prior scale mixture piMOM (spiMOM), as we introduce a hyper prior on piMOM's scale parameter. By considering an additional layer in the hierarchy, we take into account the uncertainty of τ. Consequently, spiMOM has a sub-exponential kernel, which leads to a heavier tail and more flat spikes around the origin, when compared with piMOM. Note λ has less impact than τ on determining the minimum value of $\boldsymbol{\beta}$ to be considered as non-trivial. To determine the value of λ, the Nikooienejad's procedure can be applied. Essentially, as the Nikooienejad's procedure involves random sampling among models, the robustness of the spiMOM is appealing.

Another advantage of spiMOM, as will be seen later, is its ability to detect smaller decreasing signal. Below we establish the asymptotic behavior of the posterior mode under spiMOM. We first begin with additional conditions and a lemma on the MLE.

(C1) For all $i \in \{1,...,n\}$ and $j \in \mathcal{J} \supseteq \mathcal{J}_0$, $x_{ij}\left[y_i - b'(\mathbf{x}_i^T\boldsymbol{\beta}_0)\right] = \mathcal{O}(1)$.

(C2) $\min_{j\in\mathcal{J}_0}\{\beta_{0j}\} = \mathcal{O}(n^{-m})$, $m > 0$.

Lemma 1. *Suppose for all* $\mathcal{J} \supseteq \mathcal{J}_0$ *with* $|\mathcal{J}| \leq q$ *the conditions* (A1), (A2), (B1), (B2), (C1) *and* (C2) *hold. Then* $||\hat{\boldsymbol{\beta}}_{\mathcal{J}} - \boldsymbol{\beta}_0||_2 = \mathcal{O}(n^{-1/3})$.

Remark 2. Here Lemma 1 is crucial in that it reveals the asymptotic lower bound of the posterior mode. To see this, note that the nonlocal prior is symmetric around the origin , hence resulting posterior density has posterior model (global mode) occuring at the same orthant in $\mathbb{R}^{|\mathcal{J}|}$ as the MLE, and has many other local modes in other orthants. Essentially, because of the convexity of the posterior density by which the posterior mode is contracted toward the MLE, we have the posterior mode converge toward the truth $\boldsymbol{\beta}_0$

Theorem 1. *Suppose all the conditions of Lemma 1 hold and* spiMOM *is assigned. Then, for any* $\epsilon_n^* \succ (r\lambda/n)^{1/3}$*, the posterior mode* $\hat{\boldsymbol{\beta}}_{\mathcal{J},PM}$ *satisfies*

$$p\left[\hat{\boldsymbol{\beta}}_{\mathcal{J},PM} \notin \mathcal{N}(\hat{\boldsymbol{\beta}}_{\mathcal{J}};\epsilon_n^*)\right] \to 0,$$

where $\mathcal{N}(\boldsymbol{u};\epsilon) = \left\{\boldsymbol{v} \in \mathbb{R}^{|\mathcal{J}|} : |v_j - u_j| \leq \epsilon, j \in \mathcal{J}\right\}$.

Remark 3. Theorem 1 shows that under regularity conditions, the maximum a posterior estimator derived from spiMOM is asymptotically within $(r\lambda/n)^{1/3}$-neighborhood of the MLE. While Shin et al. (2017) have proved that the maximum a posterior estimator obtained from piMOM resides at a distance of $(\tau/n)^{1/4}$ from the MLE (assuming fixed r), here the implication of Theorem 1 is apparent: the use of spiMOM improves the nonlocal prior method in detecting small decreasing coefficients.

Remark 4. It can be found in the proof of Theorem 1 that increasing r (with sample size n) though increases the penalty on complex models, but impedes the nonlocal prior method from detecting small coefficients. Therefore, to strike a balance we suggest fix r as constant.

Finally, we state below that spiMOM achieves Bayesian variable selection consistency under the proposed regularity conditions.

Corollary 1. *Suppose for all $\mathcal{J}$ with $|\mathcal{J}| \leq q$ the conditions* (A1), (A2), (B1), (B2), (C1) *and* (C2) *hold and* spiMOM *is assigned. Fix $r, \epsilon, \nu > 0$, and $0 < m < \frac{1}{3}$. Then there exist constants c_1 and c_2 such that if $c_1(1+\epsilon)(1+\nu) < \lambda^{1/6} << c_2 n^{2/9}$, then $p(\mathcal{J}_0|\mathbf{Y}_n) \xrightarrow{p} 1$.*

3 Proofs

3.1 Proof of Proposition 2

Without loss of generality, we present the proof for univariate β. The extension to multivariate $\boldsymbol{\beta}$ is straightforward due to the assumption of independence. Now with $\pi(\beta|r,\tau)$ and τ following inverse-gamma with shape $(r+1)/2$ and scale λ we can show

$$
\begin{aligned}
\pi(\beta|r,\lambda) &= \int_{\mathbb{R}^+} \pi(\beta|r,\tau)\pi(\tau|(r+1)/2,\lambda)\mathrm{d}\tau \\
&= \int_0^\infty \frac{\tau^{\frac{r}{2}}}{\Gamma(\frac{r}{2})}|\beta|^{-(r+1)} \exp\left(-\frac{\tau}{\beta^2}\right) \frac{\lambda^{\frac{r+1}{2}}}{\Gamma(\frac{r+1}{2})}\tau^{-\frac{r+1}{2}-1}\exp(-\frac{\lambda}{\tau})\mathrm{d}\tau \\
&= \frac{|\beta|^{-(r+1)}\lambda^{\frac{r+1}{2}}}{\Gamma(\frac{r}{2})\Gamma(\frac{r+1}{2})} \exp\left(-2\sqrt{\frac{\lambda}{\beta^2}}\right) \int_0^\infty 2\exp\left(-\lambda t^2 - (\sqrt{\beta^2}t)^{-2}\right)\mathrm{d}t \qquad (5) \\
&= \frac{|\beta|^{-(r+1)}\lambda^{\frac{r+1}{2}}}{\Gamma(\frac{r}{2})\Gamma(\frac{r+1}{2})} \exp\left(-2\sqrt{\frac{\lambda}{\beta^2}}\right) \int_0^\infty \frac{1}{\sqrt{\lambda}}\exp\left(-\iota^2\right)\mathrm{d}\iota \qquad (6) \\
&= \frac{|\beta|^{-(r+1)}\lambda^{\frac{r+1}{2}}\sqrt{\pi}}{2\Gamma(\frac{r}{2})\Gamma(\frac{r+1}{2})\sqrt{\lambda}} \exp\left(-2\sqrt{\frac{\lambda}{\beta^2}}\right).
\end{aligned}
$$

Note that the Equation (5) is obtained using change of variable with $\tau = t^{-2}$. The Equation (6) results from Cauchy-Schlömilch transformation

$$\int_0^\infty g\left\{\left(at - bt^{-1}\right)^2\right\} \mathrm{d}t = \frac{1}{2a}\int_0^\infty g\left(\iota^2\right) \mathrm{d}\iota.$$

3.2 Proof of Lemma 1

Since condition (C2) is satisfied, for any unit vector $\mathbf{u} \in \mathbb{R}^{|\mathcal{J}|}$, we can set $\boldsymbol{\beta}_0 = \boldsymbol{\beta} + n^{-m}\mathbf{u}$, $m > 0$. It is straightforward to see that for sufficiently large n, $\boldsymbol{\beta}$ falls into the neighborhood of $\boldsymbol{\beta}_0$ so that condition (B1) and (B2) apply. Therefore we have

$$\ell(\boldsymbol{\beta}_{\mathcal{J}}) - \ell(\boldsymbol{\beta}_0) \leq n^{-m}\mathbf{u}^T\mathcal{S}_{\mathcal{J}}(\boldsymbol{\beta}_0) - c(1-\epsilon)n^{1-2m}$$

for all $\mathcal{J} \supseteq \mathcal{J}_0$ with $|\mathcal{J}| \leq q$, which implies that $P\left\{\ell(\boldsymbol{\beta}_{\mathcal{J}}) - \ell(\boldsymbol{\beta}_0) > 0\right\} \leq \sum_{j\in\mathcal{J}} P(\mathcal{S}^2_{\mathcal{J}j}(\boldsymbol{\beta}_0) \geq n^{2-2m})$, where $\mathcal{S}_{\mathcal{J}j}$ denotes the j^{th} element of $\mathcal{S}_{\mathcal{J}}$. Note that to show the MLE is consistent with the desired convergence rate, it suffices to show that $\sum_{j\in\mathcal{J}} P(\mathcal{S}_{\mathcal{J}j}(\boldsymbol{\beta}_0) \geq n^{1-m})$ converges toward zero.

Now since the condition (C1) is satisfied, we can make use of Benette's inequality and have $P(\mathcal{S}_{\mathcal{J}j}(\boldsymbol{\beta}_0) \geq n^{1-m}) \leq \exp\left[-n^{1-2m}/(2+\mathrm{o}(1))\right]$. Observe that $|\mathcal{J}|$ is no more than $p^q \leq \exp\left\{\mathcal{O}(n^{1/3})\right\}$. Therefore, we have the MLE exist and fall within the $n^{-1/3}$-neighborhood of $\boldsymbol{\beta}_0$. The lemma is proved.

3.3 Proof of Theorem 1

First note that the log-prior density of the nonlocal prior we considered here takes the following general form $\log\pi(\boldsymbol{\beta}_{\mathcal{J}}) \propto -r\sum_{j=1}^{|\mathcal{J}|}\log(\beta^2_{\mathcal{J}j}) - \sum_{j=1}^{|\mathcal{J}|}\left\{\varphi\beta^{-2}_{\mathcal{J}j}\right\}^{\zeta}$, such that $(\zeta,\varphi) = (1,\tau)$ corresponds to the piMOM, and $(\zeta,\varphi) = (1/2,\lambda)$ corresponds to the spiMOM, respectively. Next note that, by applying second order Taylor expansion on the log-likelihood density around the MLE, we arrive at $\ell(\boldsymbol{\beta}_{\mathcal{J}}) = \ell(\hat{\boldsymbol{\beta}}_{\mathcal{J}}) - \frac{1}{2}\gamma^T\mathcal{H}_{\mathcal{J}}(\boldsymbol{\beta}^*_{\mathcal{J}})\gamma$, where $\boldsymbol{\beta}^*_{\mathcal{J}} = \hat{\boldsymbol{\beta}}_{\mathcal{J}} + \xi\gamma$, $\gamma = \boldsymbol{\beta}_{\mathcal{J}} - \hat{\boldsymbol{\beta}}_{\mathcal{J}}$, and $\xi \in [0,1]$. Putting together the log-prior density and log-likelihood, we derive the score function, i.e., the first order partial derivative of unnormalized log-posterior density with respect to $\boldsymbol{\beta}_{\mathcal{J}}$ as $\mathcal{S}^*(\boldsymbol{\beta}_{\mathcal{J}}) = -\mathcal{H}_{\mathcal{J}}(\boldsymbol{\beta}^*_{\mathcal{J}})\gamma - r\boldsymbol{\beta}^{\wedge(-1)}_{\mathcal{J}} + \varphi\zeta\boldsymbol{\beta}^{\wedge(-1-2\zeta)}_{\mathcal{J}}$, where we write $\boldsymbol{\beta}^{\wedge(-1)}_{\mathcal{J}}$ to raise each element of $\boldsymbol{\beta}_{\mathcal{J}}$ to power of minus one.

From the score function we find that, conditioning on (B1) and (B2) being true, each element of mode must satisfy $-an(\beta_j - \hat{\beta}_j) - r\beta_j^{-1} + \varphi\zeta\beta_j^{-1-2\zeta} = 0$ for some constant $a > 0$. Note here we drop the subscription of model index for simple exposition. Eventually, as

every element of mode is nontrivial, we have

$$\frac{n}{r\varphi}(\beta_j - \hat{\beta}_j)\beta_j^{1+2\zeta} \xrightarrow{p} c \tag{7}$$

for some constant c.

Without loss of generality, we consider the two modes (global and local modes) of a generic element β_j occurring given that $\text{sign}(\hat{\beta}_{j',\text{PM}}) = \text{sign}(\hat{\beta}_{j'})$ for all $j' \neq j, j' \in \mathcal{J}$. For convenient notations, we write β_{0j} for true value that corresponds to such element β_j, and write $\check{\beta}_1$ for global mode (i.e. posterior mode) of β_j whose sign is the same with its MLE, and $\check{\beta}_2$ for local mode whose sign is different from the MLE. We simply write c for generic constant if there is no confusion.

Consider first the case of fixed $\beta_{0j} \neq 0$. First note $\hat{\beta}_j \xrightarrow{p} \beta_{0j}$ implies that the global mode also converges to a constant, and thus $\check{\beta}_1^{1+2\zeta} \xrightarrow{p} c$ and $n(\check{\beta}_1 - \hat{\beta}_j)/(r\varphi) \xrightarrow{p} c$. For the local mode, since $(\check{\beta}_2 - \hat{\beta}_j) \xrightarrow{p} c$, we have $n\check{\beta}_2^{1+2\zeta}/(r\varphi) \xrightarrow{p} c$.

Now consider the case of $\beta_{0j} = 0$. First note that both modes converge toward zero at the convergence rate no faster than $n^{-1/3}$. To see this, consider $\check{\beta}_1 = \mathcal{O}(n^{-1/3-\epsilon^*})$ with some $\epsilon^* > 0$, then it follows that $n(\check{\beta}_1 - \hat{\beta}_j)\check{\beta}_1^{1+2\zeta}/(r\varphi) \asymp n^{2/3}\check{\beta}_1^{1+2\zeta}/(r\varphi)$ does not converge to some constant as required, a contradiction regardless of $\zeta = 1$ or $1/2$. In contrast, consider $\check{\beta}_1 = \mathcal{O}([r\varphi/n]^{1/3-\epsilon^*})$, then it follows that $n\check{\beta}_1^{2+2\zeta}/(r\varphi) \xrightarrow{p} c$ and therefore $\check{\beta}_1 = \mathcal{O}([r\varphi/n]^{1/4})$ if $\zeta = 1$ or $\check{\beta}_1 = \mathcal{O}([r\varphi/n]^{1/3})$ if $\zeta = 1/2$. Finally, observing that $(\check{\beta}_2 - \hat{\beta}_j) \asymp \check{\beta}_2$ (due to the opposite signs), we reach the same conclusion of the convergence rate for $\check{\beta}_2$.

Finally consider the case when $\beta_{0j} = O(n^{-m})$. We first focus on the global mode. If $\mathcal{O}(|\hat{\beta}_j - \beta_{0j}|) \preceq \mathcal{O}(n^{-m}) \prec \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$, a similar argument to the case $\beta_{0j} = 0$ implies that $\check{\beta}_1$ converges toward β_{0j} but strictly slower than $O(n^{-m})$ such that $n\check{\beta}_1^{2+2\zeta}(1 - \hat{\beta}_j/\check{\beta}_1)/(r\varphi)$ converges to a finite constant as required. However, if $\mathcal{O}(n^{-m}) \succ \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$ we shall see that assuming $|\check{\beta}_1 - \hat{\beta}_j| \succ \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$ leads to a contradiction. Note that if $|\check{\beta}_1 - \hat{\beta}_j| \succ \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$, then $|\check{\beta}_1 - \hat{\beta}_j| = \mathcal{O}(\check{\beta}_1)$ and thus $n(\check{\beta}_1 - \hat{\beta}_j)\check{\beta}_1^{1+2\zeta}/(r\varphi) = \mathcal{O}(n\check{\beta}_1^{2+2\zeta}/(r\varphi))$, which gives that $\mathcal{O}(\check{\beta}_1) \succ \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$ could not yield required results by (7), leading to a contradiction. Now consider the local mode. Similarly, if $\mathcal{O}(n^{-m}) \prec \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$, we have $\check{\beta}_2 - \hat{\beta}_j = \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$ as the signs differ, thus we have $n\,[\varphi/n]^{1/(2+2\zeta)}\,\check{\beta}_2^{1+2\zeta} \xrightarrow{p} c$. If $\mathcal{O}(n^{-m}) \succ \mathcal{O}([r\varphi/n]^{1/(2+2\zeta)})$, we have $\check{\beta}_2 - \hat{\beta}_j = \mathcal{O}(n^{-m})$ and hence we have $n^{1-m}\check{\beta}_2^{1+2\zeta} \xrightarrow{p} c$.

Overall, we have shown that in any case, there is a high probability for posterior mode to fall within the $(\varphi/n)^{1/(2+2\zeta)}$-neighborhood of the MLE. Besides, the posterior mode of spiMOM holds faster convergence rate than that of piMOM. On the other hand, the local mode also converges toward the true value, but at a rather lower rate. Finally, we observe

that increasing r or φ (with respect to n) impedes the convergence rate toward the true value. However, as φ is in the exponential kernel of piMOM (or spiMOM) while r is not, we see from Proposition 1 that a fixed r is preferred.

3.4 Proof of Corollary 1

To prove the consistency, we will show that the last two summations of (2) decrease to zero as sample size increases. Below we discuss these two cases.

Case I: $\mathcal{J} \in \mathcal{A}$

Now that spiMOM is assigned, we can rewrite the Equation (3) as

$$\log\left(\frac{\mathcal{M}_{\mathcal{J}}}{\mathcal{M}_{\mathcal{J}_0}}\right) \asymp (1+\epsilon)\log p^{(\nu+|\mathcal{J}\setminus\mathcal{J}_0|)} - \left(\sum_{j\in\mathcal{J}}\left\{\lambda\hat{\beta}^{-2}_{\mathcal{J}j,\text{PM}}\right\}^{1/2} - \sum_{j\in\mathcal{J}_0}\left\{\lambda\hat{\beta}^{-2}_{\mathcal{J}_0 j,\text{PM}}\right\}^{1/2}\right). \tag{8}$$

From the condition (A1), we have the first term in (8) such that

$$(1+\epsilon)\log p^{(\nu+|\mathcal{J}\setminus\mathcal{J}_0|)} \preceq (|\mathcal{J}\setminus\mathcal{J}_0|)(1+\epsilon)(1+\nu)n^{1/3}.$$

From the result of Theorem 1, we have the last term in (8) dominated by $\sum_{j\in\mathcal{J}\setminus\mathcal{J}_0}\left\{\varphi\hat{\beta}^{-2}_{\mathcal{J}j,\text{PM}}\right\}^{1/2}$ such that

$$\sum_{j\in\mathcal{J}\setminus\mathcal{J}_0}\left\{\varphi\hat{\beta}^{-2}_{\mathcal{J}j,\text{PM}}\right\}^{1/2} \asymp (|\mathcal{J}\setminus\mathcal{J}_0|)\mathcal{O}(\lambda^{1/6}n^{1/3}).$$

Finally, we see that if ϵ and ν are sufficiently small such that $\lambda^{1/6} > c(1+\epsilon)(1+\nu)$, then we have $\sum_{\mathcal{A}}\mathcal{M}_{\mathcal{J}}/\mathcal{M}_{\mathcal{J}_0}$ converges toward zero.

Case II: $\mathcal{J} \in \mathcal{B}$

We first cite the result of Theorem 2.2 of Barber et al. (2015) which states that

$$\ell(\hat{\boldsymbol{\beta}}_{\mathcal{J}_0}) - \ell(\hat{\boldsymbol{\beta}}_{\mathcal{J}}) \succeq cn\min_{j\in\mathcal{J}_0}|\beta_{0j}|^2 \asymp \mathcal{O}(n^{1-2m}). \tag{9}$$

Note that their result holds here because they assume more general regularity conditions. Now observe that the log-prior ratio is bounded from below by $-\sum_{j\in\mathcal{J}_0}\left\{\lambda\hat{\beta}^{-2}_{\mathcal{J}_0 j,\text{PM}}\right\}^{1/2}$. Given $m < 1/3$, we have $\min_{j\in\mathcal{J}_0}|\hat{\beta}_{0j,\text{PM}}|^2 = \mathcal{O}(n^{-m})$, and thus we have

$$-\sum_{j\in\mathcal{J}_0}\left\{\lambda\hat{\beta}^{-2}_{\mathcal{J}_0 j,\text{PM}}\right\}^{1/2} \succeq -c\lambda^{1/2}n^{-m}.$$

As a consequence, as long as $\lambda^{1/2} \prec n^{2/3}$, we have $\log(\mathcal{M}_{\mathcal{J}_0}/\mathcal{M}_{\mathcal{J}})$ dominated by $\mathcal{O}(n^{1-2m}) \succ O(n^{1/3})$ and thus arrive at the conclusion that $\sum_{\mathcal{B}}\mathcal{M}_{\mathcal{J}}/\mathcal{M}_{\mathcal{J}_0}$ converges toward zero.

4 Conclusion

In this note we discuss spiMOM, a new class of nonlocal prior that holds the robustness to specification of hyperparameters. Under certain regularity conditions, spiMOM provides maximum a posterior estimate converging at the same optimal rate as the MLE toward the truth. Overall, our approach may prove especially useful in applications of detecting small decreasing signal in high dimensional sparse data.

References

Barber, R. F., Drton, M., et al. (2015). "High-dimensional Ising model selection with Bayesian information criteria." *Electronic Journal of Statistics*, 9, 1, 567–607.

Bayarri, M. J., Berger, J. O., Forte, A., García-Donato, G., et al. (2012). "Criteria for Bayesian model choice with application to variable selection." *The Annals of statistics*, 40, 3, 1550–1577.

Chen, J. and Chen, Z. (2012). "Extended BIC for small-n-large-P sparse GLM." *Statistica Sinica*, 22, 555–574.

Johnson, V. E. and Rossell, D. (2010). "On the use of non-local prior densities in Bayesian hypothesis tests." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72, 143–170.

— (2012). "Bayesian model selection in high-dimensional settings." *Journal of the American Statistical Association*, 107, 649–660.

Liang, F., Song, Q., and Yu, K. (2013). "Bayesian subset modeling for high-dimensional generalized linear models." *Journal of the American Statistical Association*, 108, 502, 589–606.

Nikooienejad, A., Wang, W., and Johnson, V. E. (2016). "Bayesian variable selection for binary outcomes in high dimensional genomic studies using non-local priors." *Bioinformatics*, 32, 9, 1338–1345.

Shin, M., Bhattacharya, A., and Johnson, V. E. (2017). "Scalable Bayesian variable selection using nonlocal prior densities in ultrahigh-dimensional settings." *Statistica Sinica*, in press.